

Course summary:

Reinforcement learning is a powerful paradigm for learning and sequential decision making, and it is relevant to an enormous range of tasks, including robotics, game playing, consumer modeling and healthcare. The course aims to provide a hands-on introduction to the state-of-the-art techniques in reinforcement learning.

We will begin with foundations of sequential decision making and Markov Decision Processes, and quickly move on to the core challenges and recent approaches for large-scale reinforcement learning. Through a combination of lectures, assignments, readings and project, the students will become well versed in the basic RL algorithmic techniques as well as deep RL based algorithmic techniques. The assignments will involve implementing these techniques to solve [OpenAI gym environments](#) using Python ([Tensorflow](#) + numpy). We will read recently published articles on RL and utilize those ideas to potentially improve our implementations.

Coursework includes 4+ programming assignments, reading and student presentation of one recent paper, and an open-ended project. For the project, you can work with a partner. A reading list of interesting recent papers will be provided, and students can also recommend additions to the list.

The topics covered include (tentatively):

- Introduction to MDP: value-iteration, policy iteration, Q-value-iteration
- Q-learning:
 - Tabular
 - with function approximation
 - Deep Q-networks: experience replay, backpropagation
- Policy gradient methods
 - Function approximation, backpropagation
 - Trust region policy optimization
 - Actor critic methods
- Model-based RL
 - Monte Carlo Tree Search
- Further challenges
 - Exploration vs. exploitation
 - Adversarial training, Generalization
 - Multi-agent RL

Pre-requisites:

Probability and statistics, Basic optimization (e.g., familiarity with gradient descent, basic algorithm design constructs), familiarity with *Programming in python* (or experience with programming in other languages like C/C++/Matlab and willingness to learn python). Knowledge of machine learning is not required, but some basic familiarity may help.